

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau




(43) International Publication Date
9 October 2003 (09.10.2003)

PCT

(10) International Publication Number
WO 03/084106 A1

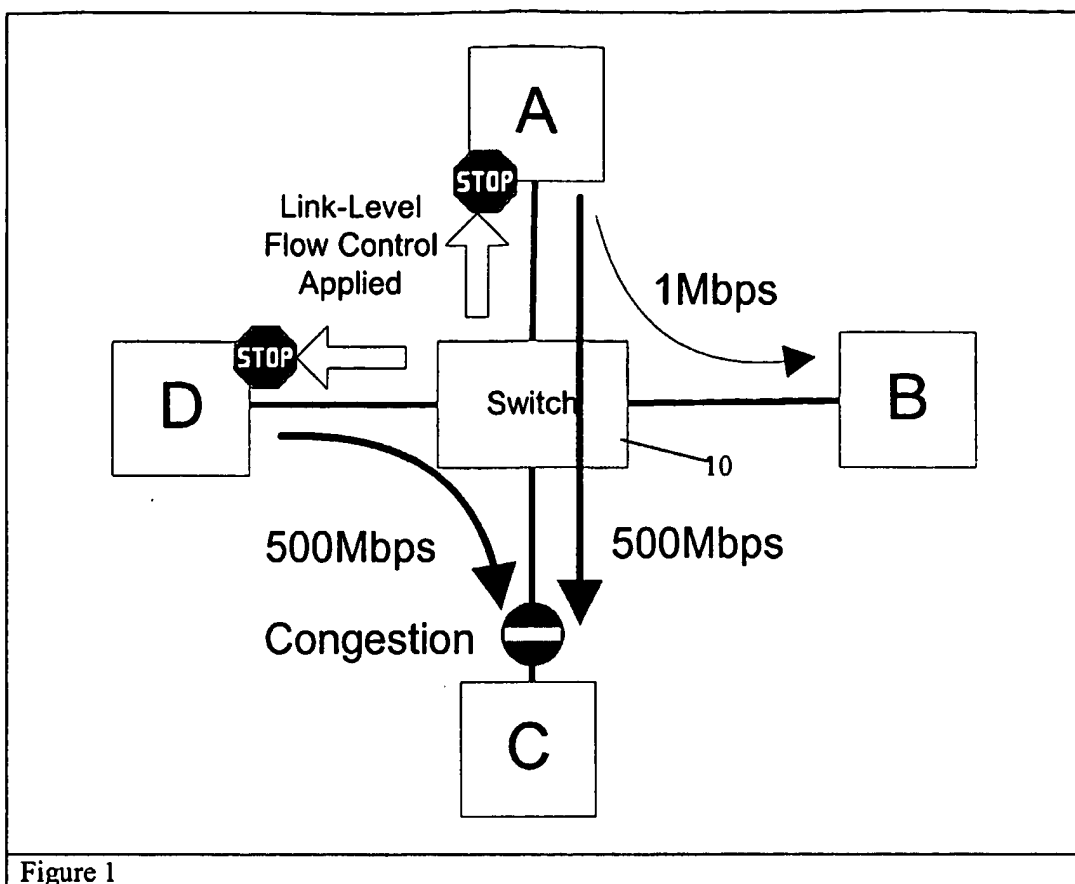
- (51) International Patent Classification⁷: **H04J 1/16**
- (21) International Application Number: PCT/US03/09920
- (22) International Filing Date: 28 March 2003 (28.03.2003)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/368,582 29 March 2002 (29.03.2002) US
Not furnished 20 March 2003 (20.03.2003) US
- (71) Applicant: **NISHAN SYSTEMS, INC.** [US/US]; 3850 North First Street, Suite B, San Jose, CA 95134-1702 (US).
- (72) Inventors: **KOAY, Englin**; 42931 Corte Caracas, Fremont, CA 94539 (US). **LATIF, Aamer**; 2047 Blue Ridge Drive, Milpitas, CA 95035 (US). **MULLENDORE, Rodney, N.**; 305 Elan Village Lane, Apt. #434, San Jose, CA 95134 (US). **WHITE, Joseph, L.**; 2000 Crystal Springs Road, #5-21, San Bruno, CA 94066 (US).
- (74) Agents: **GRAY, Gerald, T. et al.**; TOWNSEND AND TOWNSEND AND CREW LLP, Two Embarcadero Center, Eighth Floor, San Francisco, CA 94111 (US).
- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NI, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:
— with international search report
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.


(54) Title: NETWORK CONGESTION MANAGEMENT SYSTEMS AND METHODS

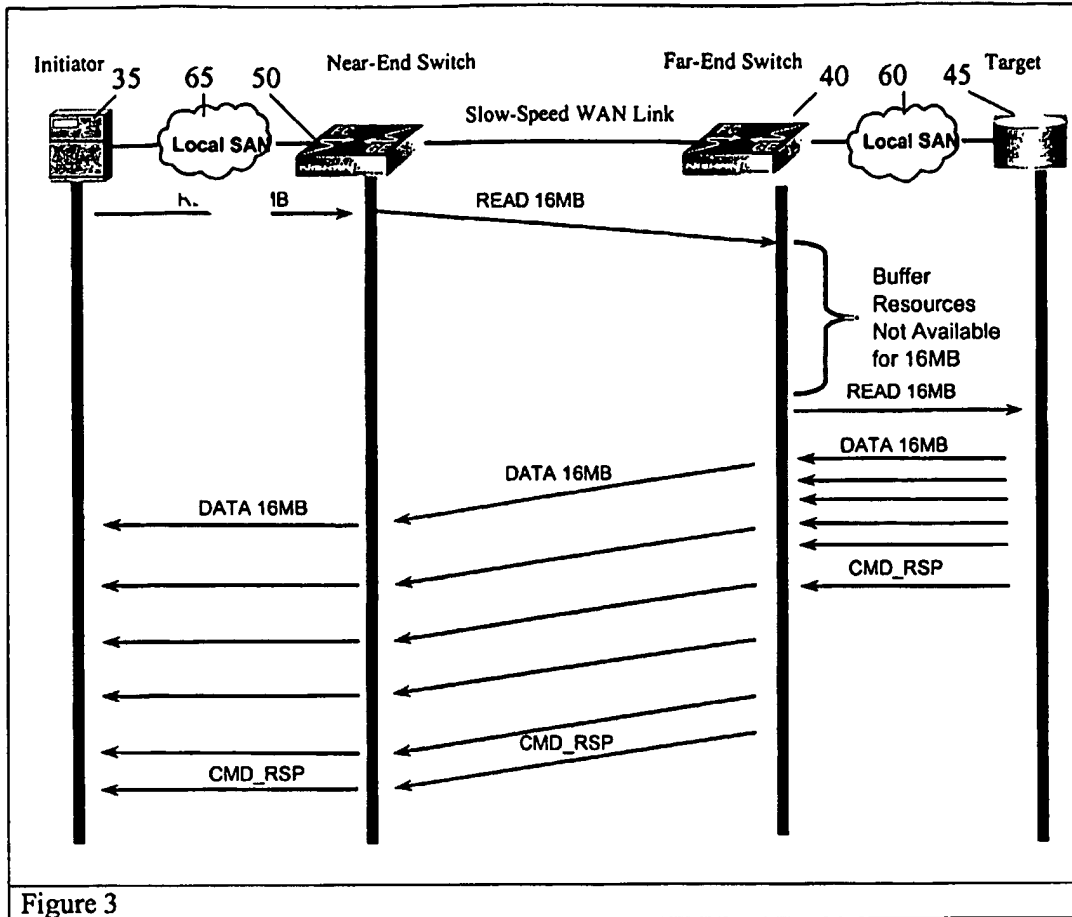


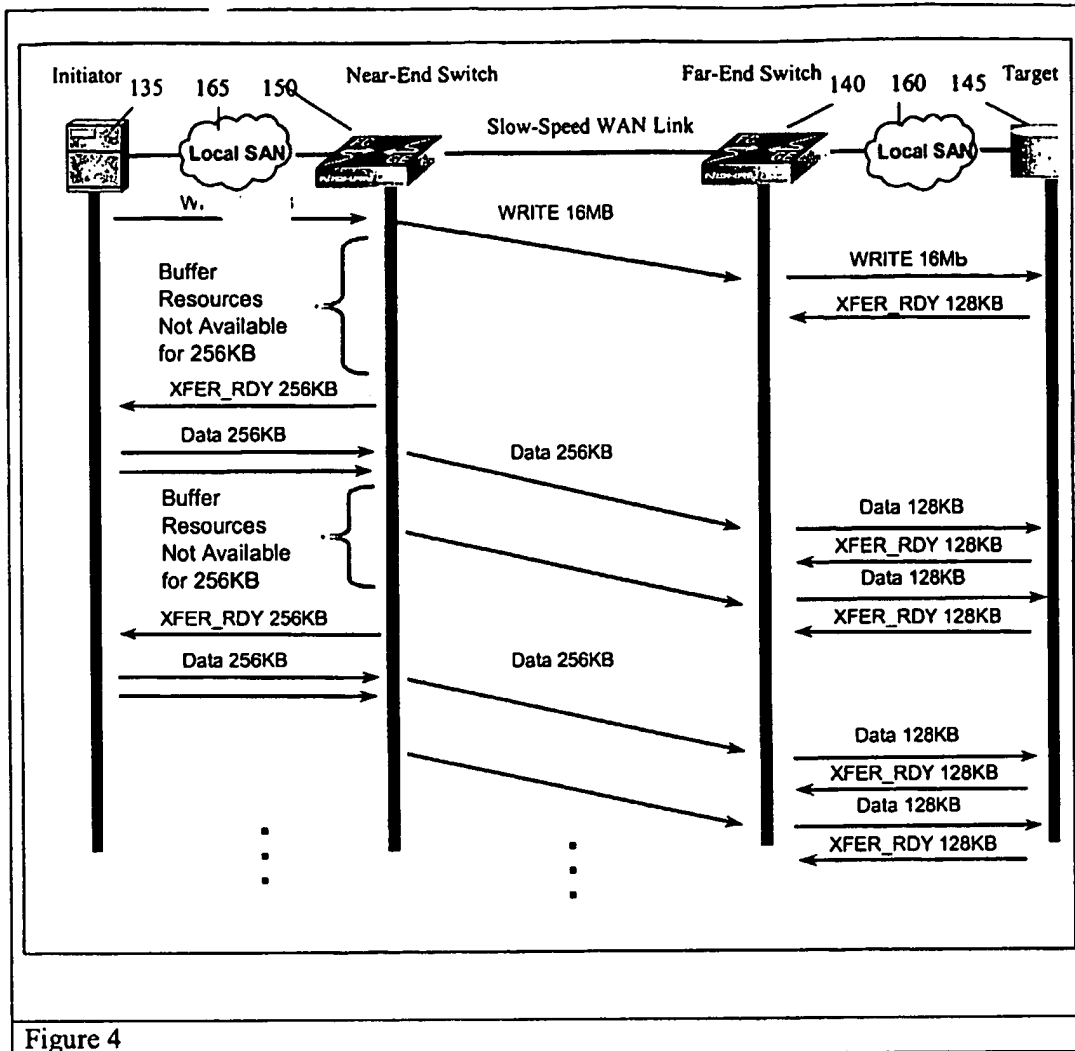
(57) Abstract: Systems, methods and software useful for overcoming network congestion problems including head-of-line blocking issues and other network congestion problems. In certain aspects, flow control mechanisms implemented in a switch device (40, 50) or other network device (40, 50) manage buffer and system level resources using a scheduler to control the amount of data requested from a local SAN fabric (60, 65). Switches (40, 50) and other network devices (40, 50) configured according to the present invention monitor each individual SCSI task, and are configured to apply as when buffering data for a slower-speed WAN link or TCP/IP based interconnects of any speed.

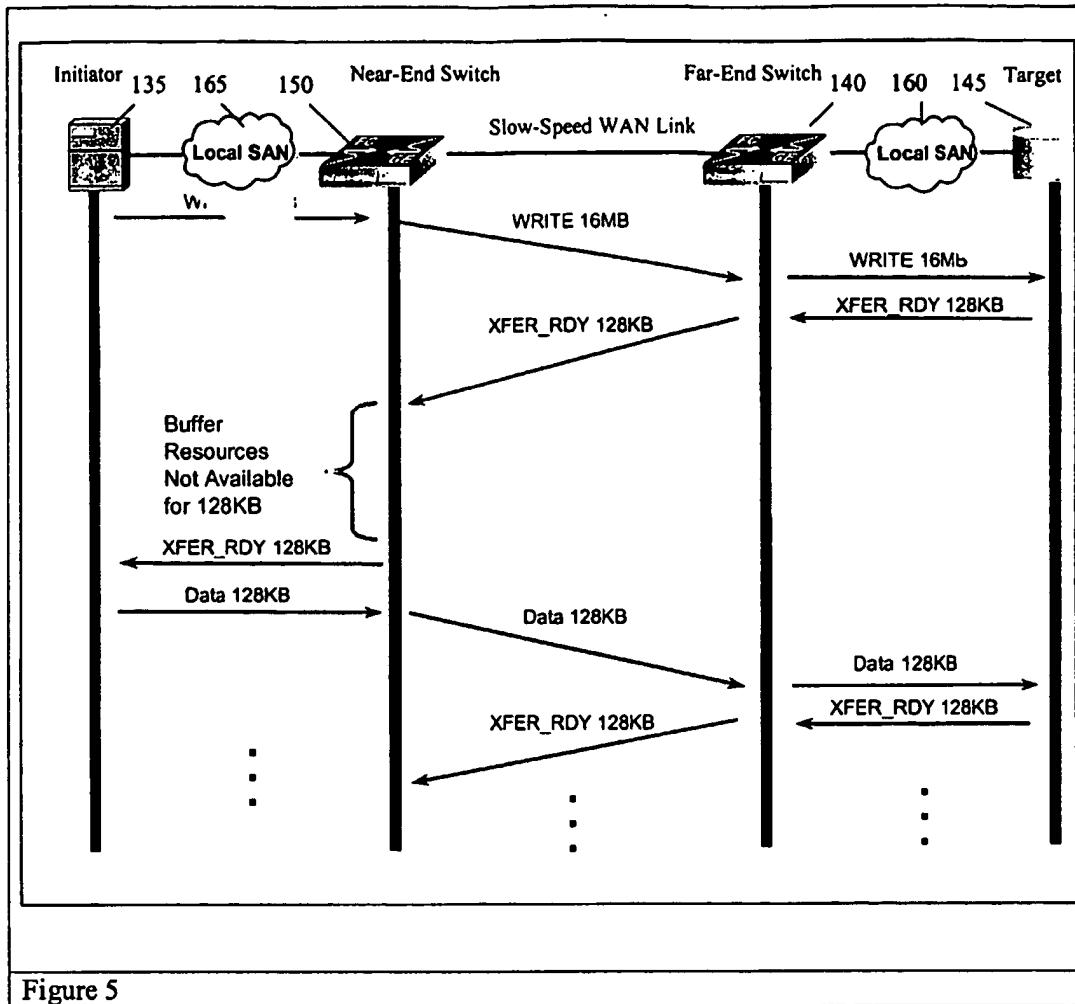
WO 03/084106 A1











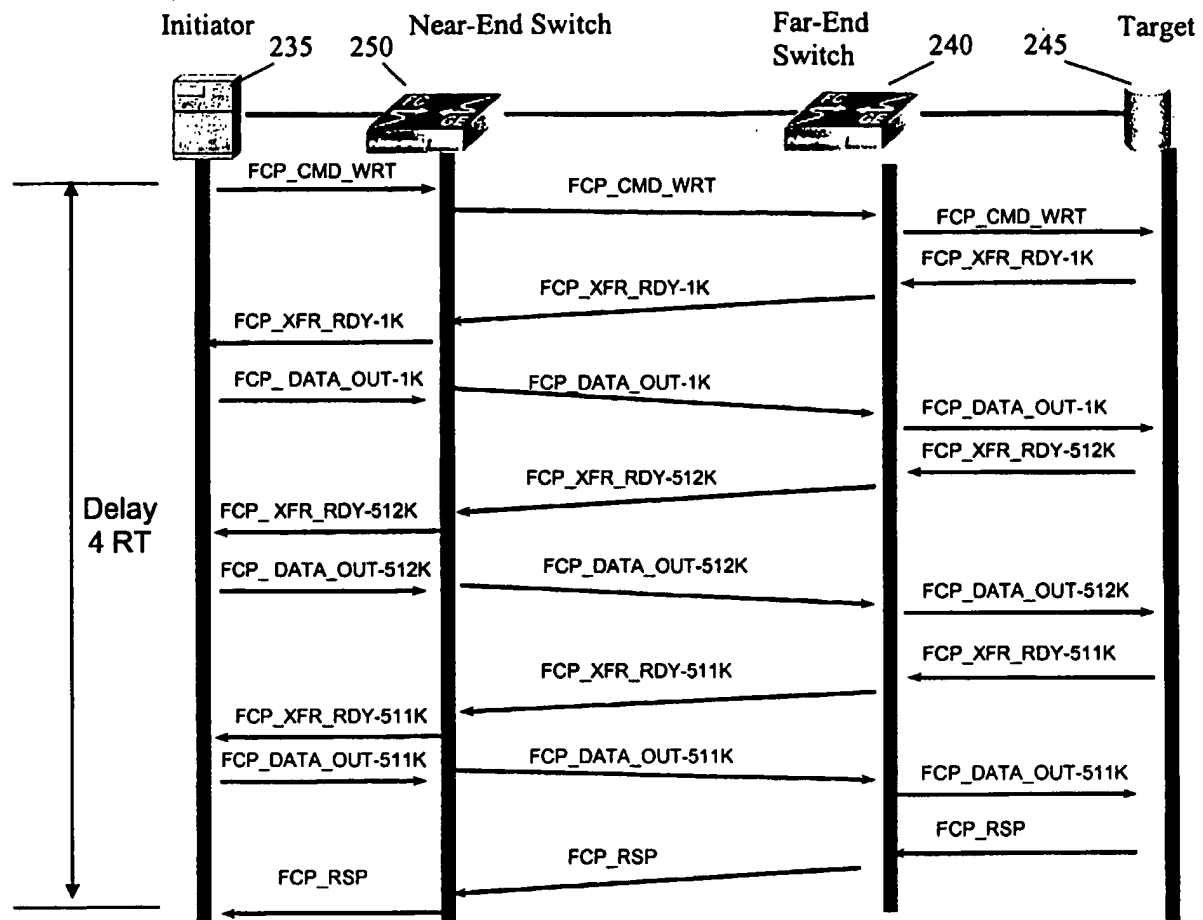


Figure 6





Figure 7



FC/Ethernet Port Configuration - 192.168.1.237

Port number: 7

Operational state: Up

Multi-function port type: Gigabit Ethernet

Port name: Port 7

Port speed (1500..1000000 Kbps): 1000000

☒ Enable port ☐ Flash LED ☒ Autonegotiations (Full duplex only)

GE Port Parameters

Port type: ☐ Layer 2 ☐ Layer 2 with iSCSI/iFCP
☐ Layer 3 ☒ Layer 3 with iSCSI/iFCP

iSCSI / iFCP Parameters

☒ Auto-reset port on severe errors ☒ iSCSI ☒ iFCP
☒ Optimize WAN throughput

	Current configuration	Active on next reset
IP address:	192.168.129.130	192.168.129.130
Subnet mask:	255.255.255.240	255.255.255.240
External router address:	192.168.129.129	192.168.129.129
Internal SAN address:	192.168.129.82	192.168.129.82
MAC address:	00:01:0F:00:53:2F	
Unicast routing protocol:	None	
Multicast routing protocol:	DVMRP	

Figure 8

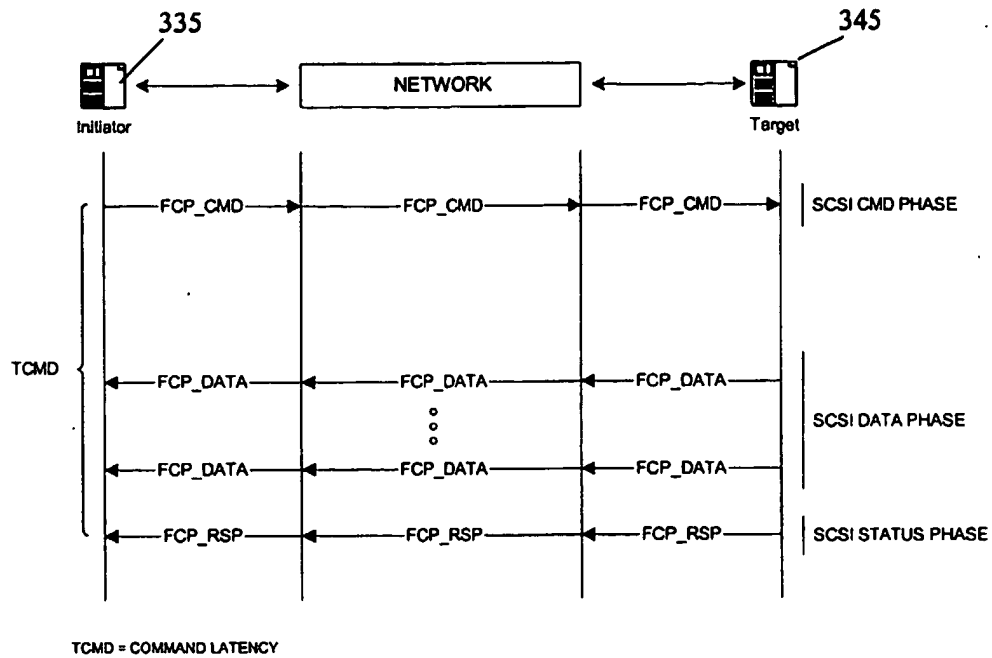


Figure 9 – SCSI Read Command In Low Latency Network




Figure 10 – SCSI Read Command In High Latency Network

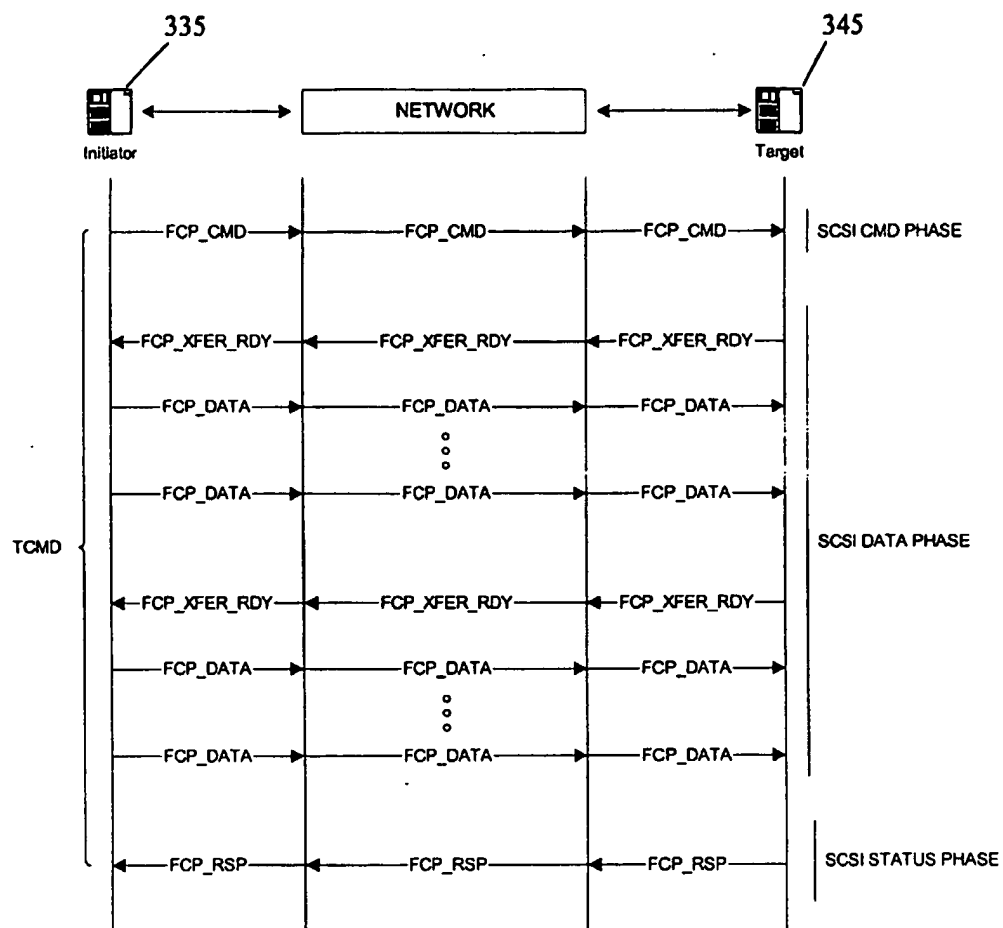
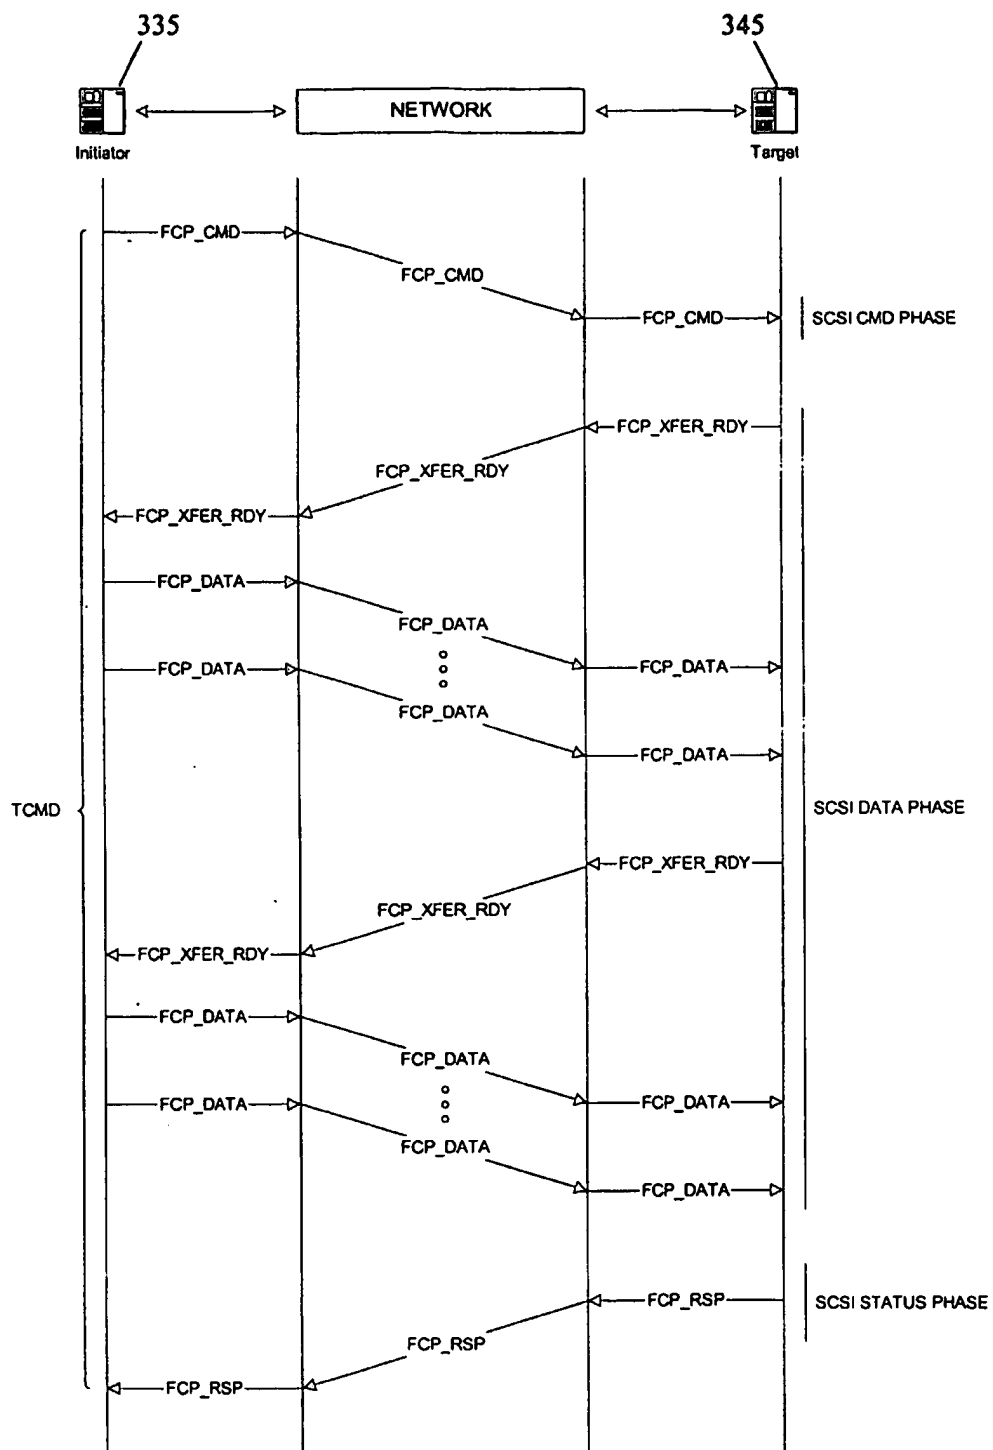


Figure 11 – SCSI Write Command In Low Latency Network

**Figure 12 – SCSI Write Command In High Latency Network**

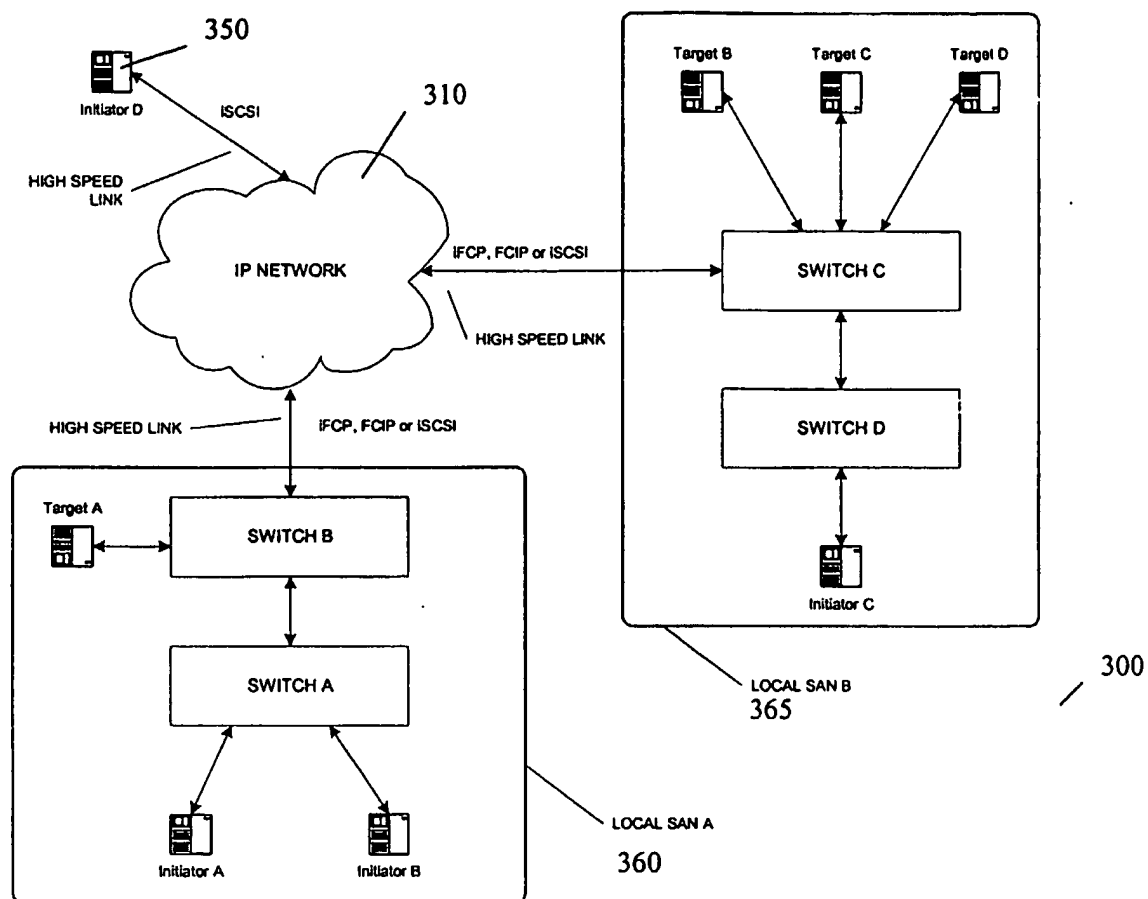


Figure 13 – Congestion Example in SAN With IP Interconnects

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US03/09920

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : H04J 1/16

US CL : 709/218, 235; 370/230, 235

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 709/218, 231, 234, 235, 250; 370/230, 235, 235.1

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EAST

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2002/0031090 A1 (KADAMBI et al.) 14 March 2002, abstract.	1-31
A	US 2001/0016878 A1 (YAMANAKA) 23 August 2001, abstract.	1-31
A	US 2001/0037406 A1 (PHILBRICK et al.) 01 November 2001, abstract.	1-31
A	US 2001/0049740 A1 (KARPOFF) 06 DECEMBER 2001, abstract.	1-31
A	US 5,592,622 A (ISFELD et al.) 07 January 1997, abstract.	1-31

☐ Further documents are listed in the continuation of Box C.
 ☐ See patent family annex.

* Special categories of cited documents:	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be of particular relevance	"X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"E" earlier document published on or after the international filing date	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"G" document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means	
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

24 JUNE 2003

Date of mailing of the international search report

11 JUL 2003

 Name and mailing address of the ISA/US
 Commissioner of Patents and Trademarks
 Box PCT
 Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

JOHN PEZZLO

Telephone No. (703) 305-3900

NETWORK CONGESTION MANAGEMENT SYSTEMS AND METHODS

CROSS-REFERENCES TO RELATED APPLICATIONS

5 This application claims the benefit of U.S. Provisional Application Serial No. 60/368,582, filed March 29, 2002, the contents of which are hereby incorporated by reference for all purposes.

BACKGROUND OF THE INVENTION

10 [01] The present invention relates generally to computer networks, and more particularly to traffic congestion management in networks.

 [02] Congestion issues are common in networks, and particularly storage networks, due to the large data flows that they must support. In Fibre Channel networks, for example,

15 congestion is typically managed through the use of link-based flow control mechanisms.

 Since there is no end-to-end flow control, head-of-line blocking of storage traffic is a common, anticipated phenomenon. Because the size of a typical Fibre Channel network is small in comparison to typical IP (Internet Protocol) networks, the impact and consequences of congestion and head-of-line blocking is limited and usually considered of minor

20 significance.

 [03] However, with the introduction of iSCSI and iFCP technologies come the potential to significantly scale the size of storage networks. Rather than the 3-4 switches typical of storage networks in the past, iSCSI and iFCP allow practically unlimited scaling in the size of storage networks. In a large IP storage network consisting of hundreds of switches, a

25 congestion issue has the potential to negatively impact the performance and reliability of a greater number of storage devices.

 [04] In addition, the use of IP introduces a greater number of link-level transports available to carry storage data, including, for example, Gigabit Ethernet, SONET, ATM, PPP, and DWDM. With the increase in types of physical transports come a much wider range of link

30 speeds at which storage data is carried, leading to potential mismatches that may compound the impact of congestion issues. In particular, congestion caused by a relatively slow link

such as a T-1 or T-3 link can cause rippling effects on the efficiency and utilization of adjacent gigabit-speed links, even if the low-speed link is rarely utilized.

[05] Head-of-line blocking is an issue for any network technology that exclusively uses link-based flow control mechanisms to manage congestion for session-based network traffic.

5 This allows for the effects of link-based flow control mechanisms, when triggered, to potentially impact sessions that are neither utilizing the congested link nor contributing to the congestion in any way.

[06] Figure 1 illustrates a basic example of head-of-line blocking in a network. As is shown, sessions to device C from devices A and D causes congestion in switch 10 when
10 devices A and D attempt to send data at a rate which is higher than device C is able to receive. Assume that the links between the switch and each device shown in Figure 1 are capable of carrying 1 Gbps (Gigabits per second) of data. Assume that devices A and D are attempting to each send 600 Mbps (megabits per second) of data to Device C. Switch 10 will be forced to buffer up some of the data since it will be receiving 1200 Mbps of data from
15 devices A and D but can only forward data at a rate of 1000 Mbps to device C. Thus, switch 10 will be accumulating data until its internal buffering is exhausted when the link-level flow control mechanism between switch 10 and devices A and D will be invoked to slow the combined data rate to 1000 Mbps. Link-level flow control thus prevents the internal buffer of switch 10 from being overflowed, thereby preventing loss of data within the network.
20 Assuming the ports are treated fairly, device A and D will each be limited to a 500 Mbps data rate. However, the link-level flow control mechanisms have no intelligence on which sessions (i.e., those directed to device C) are causing the flow control/congestion problem. Thus, other traffic that is not involved with the congestion, such as traffic from device A to device B, is affected by the link-level flow control.

25 [07] All Fibre Channel fabrics rely exclusively on the Fibre Channel link-level buffer-to-buffer credit mechanism, and are thus susceptible to head-of-line blocking issues. Until recently, Fibre Channel links were exclusively 1.0625 Gbps in throughput, and the uniformity in high-speed link throughput limited the occurrence of head-of-line blocking to those situations involving multiple session streams.

30 [08] Internet Protocol (IP) can be used to internetwork many link-level networking protocols, each characterized by different link speeds. For example, IP allows ethernet networks to be internetworked with other protocols such as ATM, Token Ring, SONET, PPP, etc... IP is "link-neutral", meaning it doesn't care what link technology is used. Due to the heterogeneity of IP transports, an end-to-end flow control mechanism such as Transmission

Control Protocol (TCP) is recommended, and a heavy reliance on link-level flow control is recognized as having unintended side-effects.

[09] The introduction of IP-based transports for connecting Fibre Channel devices or interconnecting Fibre Channel networks introduces serious congestion management issues.

5 For example, since Class 3 Fibre Channel does not have an end-to-end flow control mechanism, it must rely on link-level flow control to manage congestion and reduce packet loss. Unfortunately, this potentially raises serious head-of-line blocking issues when used with IP, since many link-level technologies used for IP are relatively slow in their link throughput compared to native Fibre Channel. Unless an end-to-end flow control mechanism
10 is introduced, a single storage session can result in serious head-of-line blocking effects that may affect traffic in the local fabric.

[10] Figure 2 illustrates the head-of-line blocking phenomenon as a result of a slow WAN link. Storage traffic simply backs up when it must ingress a slow-speed WAN link. For example, the introduction of slow speed IP links, such as T1 (1.544 Mbps), can have rippling
15 effects on congestion. As shown in Figure 2, for example, a storage session from Device A to Device D is encapsulated in IP datagrams for transmission over a slower WAN link 20. Switch #2 receives data at a faster rate than it can send over WAN link 20, and thus initiates flow control to the upstream Switch #1. Doing so, however, affects other non-related traffic that flows across the inter-switch link (ISL) between Switch #1 and Switch #2, such as traffic
20 from Device B to Device C, which competes with the WAN session traffic for available bandwidth on the ISL. Thus, the triggered link-level flow control not only impacts traffic destined for WAN link 20, but also the local high speed traffic from device B to device C. Instead, traffic from Device B to Device C will most likely have a similar throughput as that achieved over the slower WAN link 20.

25 [11] Congestion caused by head-of-line blocking may also result when Storage Networks using link level flow control are connected using high speed links such as Gigabit Ethernet or 10 Gigabit Ethernet when the protocol used on the high speed links is TCP/IP. TCP (Transmission Control Protocol) includes congestion control mechanisms as part of the protocol which dynamically change the rate at which data may be transmitted. Therefore, a
30 high speed link connected to an IP network using TCP may operate at a relatively low speed depending on the characteristics of the IP network. The data rates which can be transmitted can vary widely from the full link bandwidth (e.g. 1 Gbps or 10 Gbps) down to a few Kbps (kilobits per second).

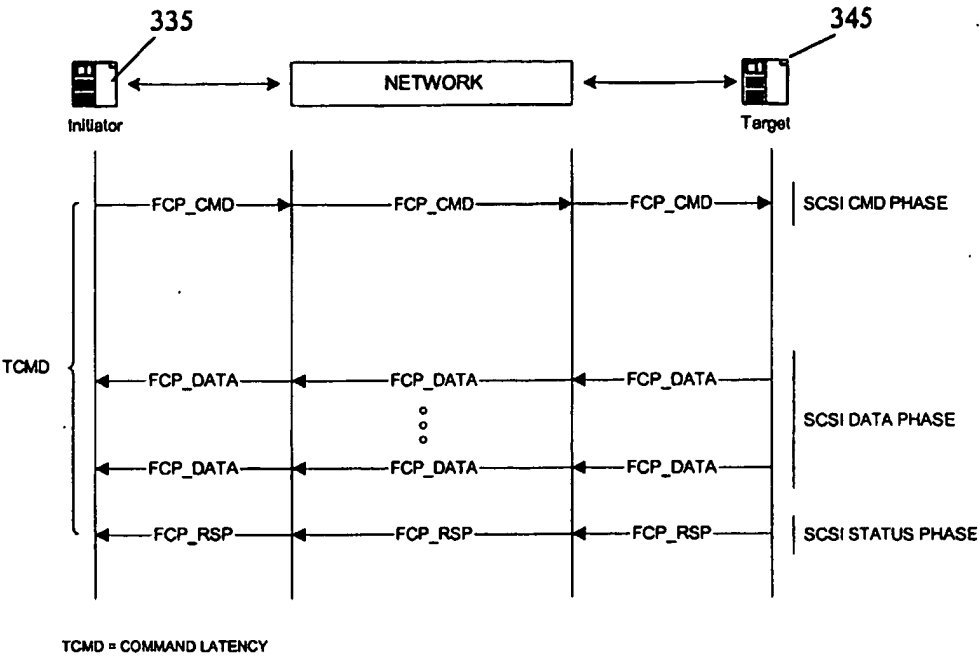


Figure 9 – SCSI Read Command In Low Latency Network




Figure 10 – SCSI Read Command In High Latency Network




Figure 11 – SCSI Write Command In Low Latency Network

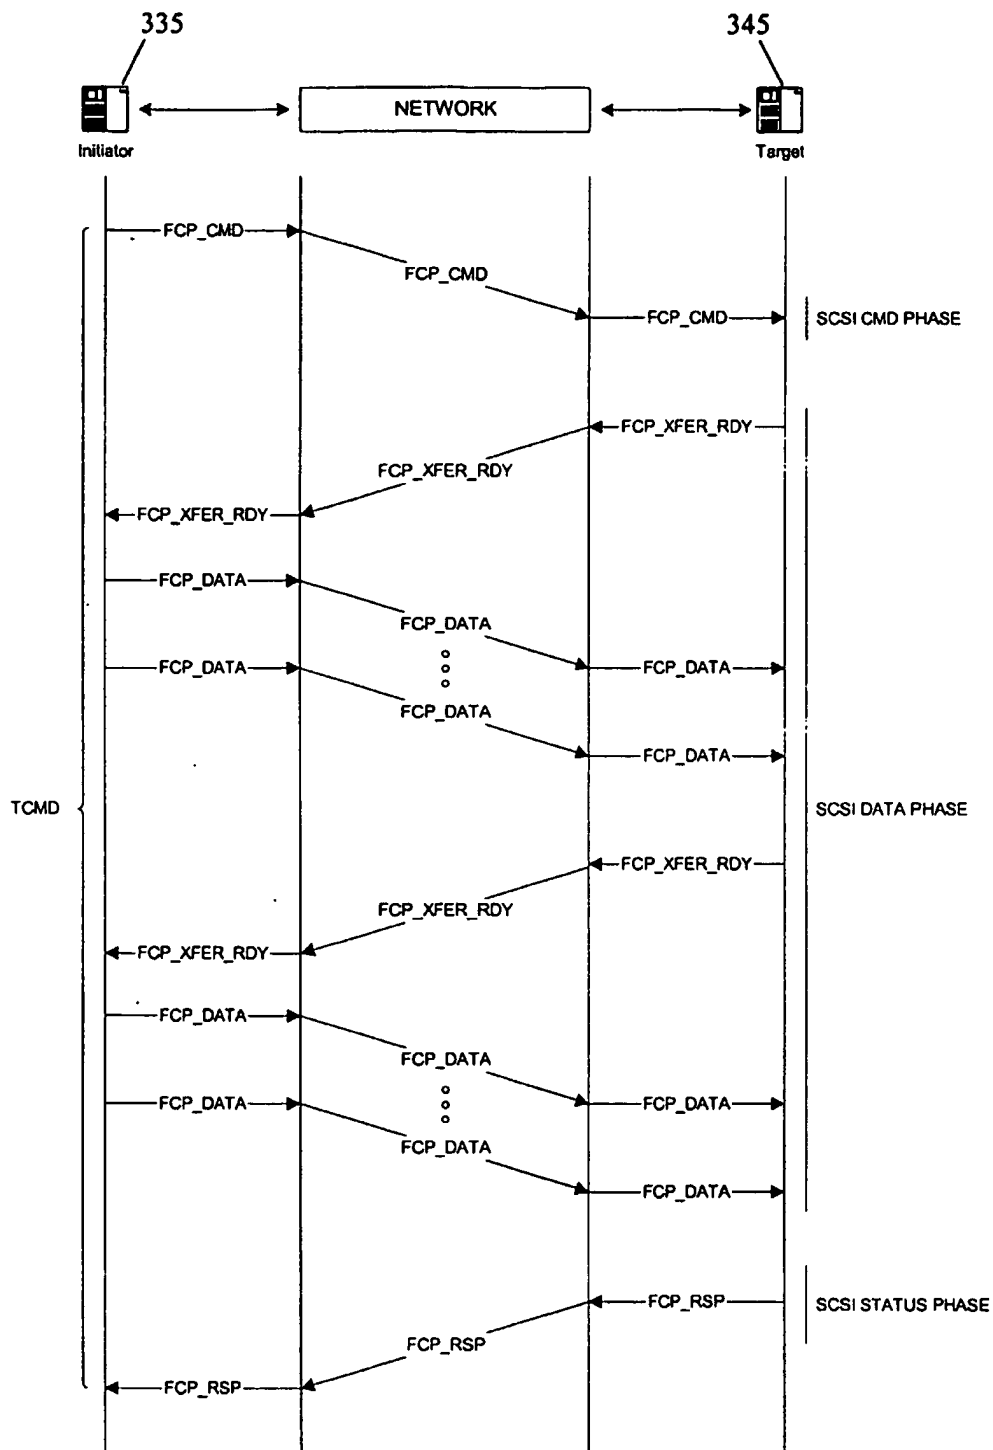


Figure 12 – SCSI Write Command In High Latency Network

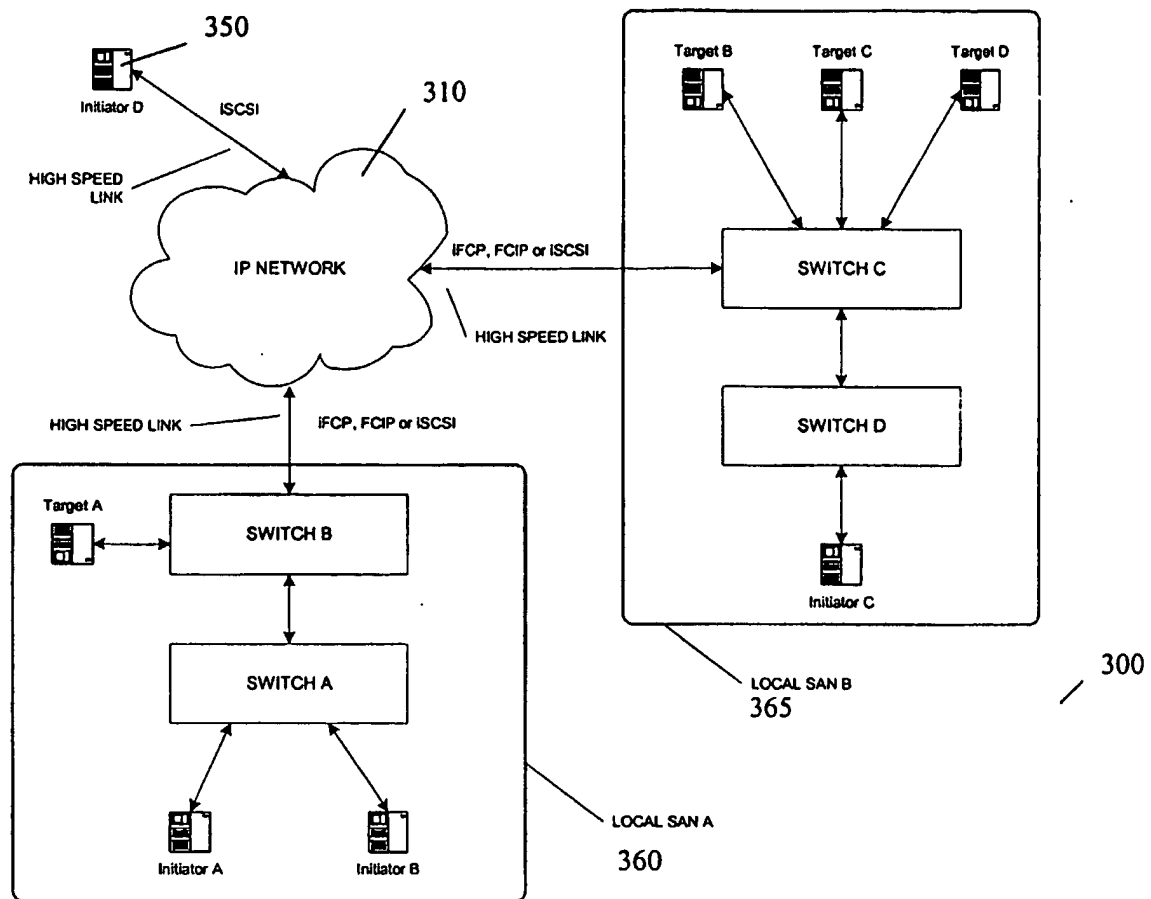


Figure 13 – Congestion Example in SAN With IP Interconnects




Figure 14 – Example Use of Rate Limiter

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US03/09920

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : H04J 1/16

US CL : 709/218, 235; 370/230, 235

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 709/218, 231, 234, 235, 250; 370/230, 235, 235.1

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EAST

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2002/0031090 A1 (KADAMBI et al.) 14 March 2002, abstract.	1-31
A	US 2001/0016878 A1 (YAMANAKA) 23 August 2001, abstract.	1-31
A	US 2001/0037406 A1 (PHILBRICK et al.) 01 November 2001, abstract.	1-31
A	US 2001/0049740 A1 (KARPOFF) 06 DECEMBER 2001, abstract.	1-31
A	US 5,592,622 A (ISFELD et al.) 07 January 1997, abstract.	1-31

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be of particular relevance	"X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"E" earlier document published on or after the international filing date	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"G" document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means	
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

24 JUNE 2003

Date of mailing of the international search report

11 JUL 2003

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

JOHN PEZZLO

Telephone No. (703) 305-3900

[12] Figure 13 shows a SAN 300 which uses an IP network 310 to interconnect Local SAN A 360, Local SAN B 365 and an initiator system or device 350 which uses iSCSI as a storage protocol. The various devices in each of the local SANs are interconnected with a link level flow control based network such as Fibre Channel. The protocols used to interconnect the local SANs to the IP network 310 may be any TCP/IP based storage protocol such as, for example, iFCP, FCIP and iSCSI. The same congestion problems which can occur when 2 local SAN networks are interconnected with a slow speed WAN link can also occur in the SAN 300 shown in Figure 13 because the high speed links connecting switches B and C to the IP network 310 may have only a fraction of their bandwidth utilized. The usable bandwidth on the high speed links may be limited by the TCP/IP protocol which reduces the bandwidth used when it detects frames being discarded in the IP network. The used bandwidth will gradually be increased until frames are once again dropped. However, the recovery from dropped frames by TCP often results in no data being transmitted for periods of 1 or more seconds.

[13] The introduction of data links that have a high latency, such as IP-based WAN links, can result in a significant degradation in write performance. Read performance can also be negatively impacted but typically to a lesser extent than write performance. The drop in performance is typically due to handshaking within the protocol used to carry the SCSI commands. Figure 9 shows an example of a SCSI read command using FCP (Fibre Channel Protocol) in a low latency network. In this example, an initiator 335 issues a read command (FCP_CMD) to a target device 345 requesting that the target return a specific group of data. The target returns the requested data in (FCP_DATA) packets on the network followed by the command status (e.g., in an FCP_RSP frame). In a low latency network, the time required for the read command to complete is dominated by the time required by the target to process the command, retrieve the requested data from memory media (e.g., disk drives, magnetic tape, etc.) and transmit the data and status in packets to the initiator. The addition of latency in the network increases the time required to complete a read command by the total round trip time (RT) of the network (a network with a latency of 5 ms in each direction has an RT of 10 ms). Figure 10 shows how the read command is affected by network latency in a high latency network. The time required for the read command to complete is increased by RT. For example, if a read command would normally complete in 10 ms (milliseconds) in a network with no latency, the same command would require 60 ms to complete in a network with 50 ms of latency.

[14] Figure 11 shows an example of a SCSI write command performed using FCP in a low latency network. The initiator 335 issues the write command request (FCP_CMD). The target 345 receives the write command and returns an FCP_XFER_RDY frame when it is ready to accept the data for the write command. The target indicates in the FCP_XFER_RDY frame the amount of the command write data it is requesting from the initiator. The target may request any amount of the data. The initiator sends the requested data to the target in FCP_DATA frames. When the target receives all of the requested data, it either requests additional data by issuing another FCP_XFER_RDY frame if all of the data has not yet been received or returns the SCSI status in an FCP_RSP frame completing the SCSI command. For example, if the initiator issued a 256 KB (kilobyte) write command, the target could return a XFER_RDY frame requesting 64 KB of the data. When this data was transferred, the target could request another 64 KB, then another 64 KB until all of the data was transferred. When all of the data has been transferred, the target would reply with an FCP_RSP frame indicating the status for the SCSI command. Alternatively, the target could have issued a single request for 256 KB of data or for any combination of requests which summed to 256 KB. Note that the target can not issue another request for data until it has received the data from an earlier data request.

[15] Figure 12 shows the effect of performing write commands over a network with high latency. The network latency has a greater effect on the time required to complete write commands than for read commands due to the additional handshakes between the initiator 335 and target 345. The write command completion time will be delayed an additional $N \times RT$ where N = the number of FCP_XFER_RDY+1 issued by the target. For example, assume a write command would complete in 10 ms in a network with no latency. The same command on a network with a 50 ms RT would require 110 ms to complete if the target issues a single FCP_XFER_RDY frame. The required time could be much higher if the target issues multiple FCP_XFER_RDY frames. For example, if the target issued 4 XFER_RDY frames, the delay would increase to 260 ms.

[16] It is therefore desirable to provide congestion management systems, methods and software that avoid or significantly reduce the effects of head-of-line blocking and network latency. Such technologies should allow for the full and efficient utilization of slow-speed and/or high latency links within a network, e.g., storage area network (SAN), without impacting the performance of upstream high-speed links.

BRIEF SUMMARY OF THE INVENTION

[17] The present invention provides systems, methods and software useful for overcoming network congestion problems including head-of-line blocking issues and other network congestion problems. In certain aspects, the present invention is particularly applicable for use with any networked transport mechanisms used to carry SCSI operations between SCSI initiator and target devices, including TCP/IP for example.

[18] According to the present invention, congestion management systems and methods are provided to overcome head-of-line blocking issues resulting from slower-speed network links, such as low speed WAN links or links using a TCP/IP based storage protocol. The flow-control congestion management systems and methods of the present invention advantageously prevent head-of-line blocking in each local SAN fabric. According to one aspect, such flow control mechanisms manage buffer and system level resources on a per-task basis. According to another aspect, such flow control mechanisms manage buffer and system level resources using a scheduler to control the amount of data requested from the local SAN fabric. Switches and other network devices configured according to the present invention monitor each individual SCSI task, and are capable of applying flow control measures to each active session when buffering resources become scarce, such as when buffering data for a slower-speed WAN link or TCP/IP based interconnects of any speed.

[19] A congestion management system, or Congestion Manager, of the present invention is a valuable component of an integrated storage network that links local SAN fabrics (implemented with a link level flow control protocol) over a wide geographic distance. The Congestion Manager advantageously allows local SANs to function independently, without being adversely impacted by head-of-line blocking, for example, when they are connected to remote SAN fabrics using long-distance WAN links or TCP/IP. Switches configured with a Congestion Manager according to the present invention can use appropriate end-to-end flow control in an appropriate manner that minimizes disruption in the local high-performance SAN.

[20] According to one aspect of the present invention, a method is provided for reducing network congestion. The method typically includes receiving a message by a network device coupling a high speed network link with a low speed or TCP/IP based network link, wherein the network device has a buffer memory, and wherein the message is sent from a requesting device to a destination device requesting that data be sent over the low speed or TCP/IP based link from the destination device to the requesting device. The method also typically includes determining whether the buffer memory has sufficient space to buffer the amount of data

identified by the message request. If the buffer has sufficient space, the method typically includes transferring the message to the destination device and buffering the requested data received from the destination device in response to the message, wherein the requested data is sent over the low speed or TCP/IP based link destined for the requesting device. If the buffer
5 doesn't have sufficient space, the method typically includes holding the message until the buffer has sufficient space.

[21] According to another aspect of the present invention, a method is provided for reducing network congestion. The method typically includes monitoring operation requests received by a network device coupling one or more high speed network links with a low
10 speed or TCP/IP based network link, wherein the network device has a buffer memory, and wherein the requests are sent between requesting devices and destination devices identifying data to be sent over the low speed or TCP/IP based link. For each received operation request, the method typically includes determining whether the buffer memory has sufficient space to buffer the amount of data identified by the request, and if so, transferring the operation
15 request to the destination device and buffering the identified data received from the destination device, wherein the requested data is sent over the low speed or TCP/IP based link destined for the requesting device, and if not, holding the operation request until the buffer has sufficient space.

[22] According to another aspect of the present invention, a method is provided for
20 reducing network congestion. The method typically includes monitoring operation requests received by a network device coupling one or more high speed network links with a low speed or TCP/IP based network link, wherein the network device has a buffer memory, and wherein the requests are sent between requesting and destination devices identifying data to be sent over the low speed or TCP/IP based link. For each received operation request, the
25 method typically includes controlling the rate at which the received operation requests are forwarded based on the amount of data to be returned. In one aspect, the received operation requests are forwarded by the network device such that the rate of requested data returned is substantially equal to the rate of the low speed or TCP/IP based network link. In certain aspects, forwarding of operation requests ceases temporarily when a threshold on the amount
30 of available buffer memory remaining (e.g. 5%, 10%, 20%, etc.) in the network device is exceeded.

[23] According to a further aspect of the present invention, a method is provided for enhancing write performance in a network including first and second switch devices coupled over a low speed or TCP/IP based network link, wherein the first switch device is coupled to

an initiator device over a first high speed network link, and wherein the second switch device is coupled to a target device over a second high speed network link. The method typically includes automatically responding to a write request received by the first switch from the initiator with one or more ready-to-transfer messages on behalf of the target device, the
5 ready-to-transfer messages requesting the write data from the initiator, and sending the write request to the target via the second switch device. The method also typically includes receiving the write data from the initiator, the write data being sent in response to the ready-to-transfer messages, and automatically sending the write data from the first switch to the second switch over the low speed or TCP/IP based network link so that the write data is
10 stored on the second switch device. When the target sends one or more ready-to-transfer messages requesting all or a portion of the write data, the second switch is able to immediately respond with the requested amount of the stored write data.

[24] According to yet a further aspect of the present invention, a method is provided for enhancing write performance in a network including first and second switch devices coupled
15 over a first network link, wherein the first switch device is coupled to an initiator device over a second network link, and wherein the second switch device is coupled to a target device over a third network link, wherein the first network link has a high latency. The method typically includes automatically responding to a write request received by the first switch from the initiator with one or more transfer messages on behalf of the target device, the
20 transfer messages requesting the write data from the initiator, and sending the write request to the target via the second switch device. The method also typically includes receiving the write data from the initiator, the write data being sent in response to the transfer message, and automatically sending the write data from the first switch to the second switch over the first network link so that the write data is stored on the second switch device. When the target
25 sends one or more transfer messages requesting all or a portion of the write data, the second switch is able to immediately respond with the requested amount of the stored write data.

[25] According to yet another aspect of the present invention, a network switch device is provided. The Switch typically includes a first port for coupling to a high speed network link, a second port for coupling to a low speed or TCP/IP based network link, a buffer
30 memory, and a congestion management module executing on the switch device. The module is typically configured to monitor messages being sent between requesting and destination devices requesting that data be sent over the low speed or TCP/IP based link, and to determine, for each message, whether the buffer memory has sufficient space to buffer the amount of data identified by the message. If it is determined that there is sufficient space, the

switch device transfers the message to the destination device and buffers the requested data received from the destination device in response to the message, wherein the requested data is sent over the low speed or TCP/IP based link destined for the requesting device, and if there is not sufficient space, the switch device holds the message until the buffer has sufficient

5 space.

[26] According to still a further aspect of the present invention, a control module executing on a network switch device is provided. The switch device typically includes a processor, a buffer, a first port for coupling to a low speed or TCP/IP based network link, and one or more second ports for coupling to one or more high speed network links. The module is typically
10 configured with instructions to monitor messages being sent between requesting and destination devices requesting that data be sent over the low speed or TCP/IP based link, and to determine, for each message, whether the buffer memory has sufficient space to buffer the amount of data identified by the message. The module is also typically configured to control the switch device, if it is determined that there is sufficient space, to transfer the message to
15 the destination device and buffer the requested data received from the destination device in response to the message, wherein the requested data is sent over the low speed or TCP/IP link destined for the requesting device, and if there is not sufficient space, to control the switch device to hold the message until the buffer has sufficient space.

[27] According to still a further aspect of the present invention, a control module executing
20 on a network switch device is provided. The switch device typically includes a processor, a buffer, a first port for coupling to a low speed or TCP/IP based network link, and one or more second ports for coupling to one or more high speed network links. The module is typically configured with instructions to monitor messages being sent between requesting and destination devices requesting that data be sent over the low speed or TCP/IP based link. The
25 module is also typically configured to implement a rate limiting function to determine if a message should be forwarded to the destination device, wherein the module controls the switch device to transfer the message to the destination device and buffer the requested data received from the device in response to the transferred message if it is determined that a rate limit is not violated, wherein the requested data is sent over the low speed or TCP/IP link
30 destined for the requesting device, and wherein the module controls the switch device to hold the message until the rate limit is no longer violated if sending the message would violate the rate limit.

[28] According to still another aspect of the present invention, a write enhancement module executing on a network switch device is provided. The switch device typically

includes a processor, a buffer, a first port for coupling to a low speed or high latency network link, and a second port for coupling to a high speed network link. The module is typically configured with instructions to control the switch to automatically respond to a write request received by the switch over the high speed network link from an initiator device with one or more transfer messages on behalf of the target device, the transfer messages requesting the write data from the initiator, and to send the write request to the target via a second switch device over the low speed or high latency network link. The module is also typically configured to control the switch to receive the write data from the initiator, said write data being sent in response to the transfer messages, and automatically send the write data to the second switch over the low speed or high latency network link so that the write data is stored on the second switch device, such that when the target sends one or more transfer messages requesting all or a portion of the write data, the second switch is able to immediately respond with the requested amount of the stored write data.

[29] Reference to the remaining portions of the specification, including the drawings and claims, will realize other features and advantages of the present invention. Further features and advantages of the present invention, as well as the structure and operation of various embodiments of the present invention, are described in detail below with respect to the accompanying drawings. In the drawings, like reference numbers indicate identical or functionally similar elements.

BRIEF DESCRIPTION OF THE DRAWINGS

- [30] Figure 1 illustrates how multiple sessions to Device C cause congestion, triggering link flow control to device A.
- [31] Figure 2 illustrates how the introduction of slow-speed IP links, such as a T1 link, may have rippling effects on congestion.
- [32] Figure 3 illustrates end-to-end flow control applied to a read operation according to an embodiment of the present invention.
- [33] Figure 4 illustrates flow control of write operations according to one embodiment of the present invention.
- [34] Figure 5 illustrates an embodiment of the present invention with Fast Write disabled.
- [35] Figure 6 shows a standard write transaction for a 1MB write transaction.
- [36] Figure 7 illustrates a near-end switch operating in Fast Write mode according to one embodiment.

[37] Figure 8 illustrates a sample user interface screen provided by software executing in a computer system coupled to a network according to one embodiment.

[38] Figure 9 shows an example of a SCSI read command using FCP (Fibre Channel Protocol) in a low latency network.

5 [39] Figure 10 shows how a read command is affected by network latency in a high latency network.

[40] Figure 11 shows an example of a SCSI write command performed using FCP in a low latency network.

10 [41] Figure 12 shows the effect of performing write commands over a network with high latency.

[42] Figure 13 shows a SAN which uses an IP network to interconnect Local SAN A, Local SAN B and an initiator device which uses iSCSI as the storage protocol.

[43] Figure 14 shows an example use of a rate limiter module in a switch B to improve the congestion behavior of the local SAN.

15

DETAILED DESCRIPTION OF THE INVENTION

[44] A network device, such as a switch device or other device, is configured with a congestion management module ("Congestion Manager") according to one embodiment of the present invention. The Congestion Manager, in certain aspects, is configured to monitor traffic sessions flowing through the device and to implement resource management algorithms responsive to detected traffic. In certain aspects, the Congestion Manager monitors each SCSI task and implements an intelligent algorithm to ensure an optimal dynamic allocation of finite buffer resources to each task. For example, to overcome head-of-line blocking issues resulting from slower-speed links, such as low speed WAN links or TCP/IP based links, SCSI-level end-to-end flow control congestion management is provided by a Congestion Manager in a switch device to advantageously prevent head-of-line blocking in each local SAN fabric. Such flow control mechanisms manage buffer and system level resources on a per-task basis. Switches and other network devices configured with a Congestion Manager according to the present invention preferably monitor each individual SCSI task, and are capable of applying flow control measures to each active session when buffering resources become scarce, such as when buffering data for a slower-speed WAN link or TCP/IP based link.

20

25

30

[45] Figure 3 illustrates flow control measures implemented by a Congestion Manager implemented in a switch device 40 according to one embodiment. When the Congestion Manager detects a command (e.g., either a read command or Ready-To-Transfer (RTT) message) for a SCSI operation that would consume excessive switch buffer resources, that command is held by the switch until sufficient buffer resources become available to receive the next complete data transfer segment for that command. This action serves to delay progress for that SCSI operation, providing time for resources on the switch to become available. If sufficient buffer resources are currently available, the command is not held. Once sufficient buffer resources are available, the command is released, allowing the SCSI operation to generate an amount of data that matches available resources.

[46] As shown in Figure 3, for example, when the Congestion Manager resident on switch 40 ("Far-End Switch" relative to initiator 35) detects a read command issued by initiator 35, the read operations are held until buffer resources on switch 40 are sufficient to receive all of the data requested from target 45 for the read command. For read commands, the Congestion Manager operates in the same manner regardless of whether the Fast Write capability (discussed below) is utilized or not. The switch 40 remote to the SCSI initiator 35 (closest to the target 45) monitors SCSI READ commands traveling to local targets. When switch 40 runs low on buffer resources needed to receive data from the target(s) for the next arriving SCSI READ command, the Congestion Manager instructs the switch 40 to hold the READ command and all others that follow until additional buffer resources become available. When additional buffer resources are freed, the READ command is released and allowed to travel to the target, causing the target to generate the READ data.

[47] In one embodiment, the Congestion Manager includes a rate limiter module configured to control the rate at which data requests are issued into a local SAN to aid in minimizing congestion. Figure 14 shows an example use of a rate limiter module implemented in switch B 440 to improve the congestion behavior of the local SAN 460. In this example, assume that each local SAN 460 and 465 uses 1 Gbps links and the IP network 430 supports a data rate of 100 Mbps. The initiator 435 issues 8 read commands, each requesting 1 MB of data, which are transferred to Switch B. Switch B would normally immediately forward the individual commands onto the target. However, in one embodiment, the rate limiter module is configured to add delay between the commands to match the rate of the IP network. In this example, a 1 MB read request represents 8 Mbps of bandwidth or 80 ms of time on the 100 Mbps IP network, therefore, the rate limiter inserts 80 ms of time between the forwarded read requests in one embodiment. The use of a rate limiter

advantageously makes the data traffic in the local SAN 460 less bursty. In the example of Figure 14, the local SAN 460 would experience smaller bursts (1 MB) of traffic spaced at even intervals (80 ms) as opposed to a single large burst (8 MB).

[48] In another embodiment, the rate limiter module is configured to add delay resulting in a data transfer rate to the network device from the target(s) that is less than or equal to the rate of the IP network. Alternately, or additionally, the rate limiter is configured in one embodiment to monitor the buffer resources on the network device, and if the amount of available memory resources currently available is less than a threshold amount (e.g., 5%, 10%, 20%, 30%, etc.) to temporarily hold the commands until sufficient memory resources become available.

[49] Referring back to Figure 3, without a Congestion Manager, SCSI operations would pass through the network device (e.g., switch 40) regardless of available resources to handle data generated by the operations. If the operations generated a greater amount of data than was available in the local device buffers, flow control would be applied at the link level, preventing upstream nodes (e.g., devices within local SAN 60, such as target 45, a storage device controller, switch device, etc.) from overflowing the available buffers. As a result of the link-level flow control, head-of-line blocking is a possibility for traffic utilizing the upstream device/node.

[50] In addition, a Congestion Manager advantageously helps reduce I/O latency. Consider, for example, the case of the initiator 35 issuing both read and write commands. Without a Congestion Manager resident on switch 40, the read commands would be delayed behind the write data that is queued. If there is less write data queued, the read commands, when issued by the initiator, will propagate to the target with less delay. Likewise, the RTT issued by the target will also propagate more quickly to the initiator (or switch 50) when the amount of read data outstanding is limited.

[51] In certain aspects, once buffer resources have been committed to a given SCSI operation, those resources are reserved for a finite period until either 1) data from that operation is received and the committed buffer resources are utilized, or 2) the finite period expires and the committed resources are reallocated to a different SCSI operation.

[52] In certain aspects, the Congestion Manager is configured to monitor all active SCSI tasks and to allocate available buffer resources in a manner that ensures fairness among the active tasks. For example, if one particular SCSI task consumes a disproportionate amount of buffer space, then in order to ensure fairness, the Congestion Manager holds and delays the SCSI messages for that task in order to provide more buffer resources to other SCSI tasks.

Within each SCSI task, SCSI messages are preferably held and released on a First-In-First-Out (FIFO) basis to prevent re-ordering of messages within the task.

[53] In one embodiment, if the SCSI transport protocol used provides information that identifies the SCSI initiator device for each SCSI task, the Congestion Manager allocates resources fairly among all known SCSI initiator devices. According to one embodiment, for example, the Congestion manager includes a scheduling algorithm module configured to implement a scheduling algorithm, such as a weighted fair (equal) queuing algorithm, for allocating memory to outstanding tasks. Other scheduling algorithms may be used such as a round robin or strict priority algorithm.

[54] In certain aspects, the Congestion Manager is configured to monitor different types of SCSI messages, depending on the direction of the operation. For example, in one embodiment, for read operations, the Congestion Manager monitors SCSI READ commands entering the locally-attached SAN fabric, and for write operations, the Congestion Manager monitors SCSI Ready-To-Transfer (RTT) messages leaving the locally-attached SAN fabric and entering the long distance WAN link.

[55] *Congestion Manager for Write Operations*

[56] The process to control the flow of write operations is similar to that described above for read operations, except the Congestion Manager monitors SCSI Ready-to-Transfer (RTT) messages. Because data flows in the opposite direction from read operations, the switch local to the SCSI initiator applies the flow control measures on the RTT messages. Thus, the Congestion Manager implemented in switch 50 of Figure 3 monitors RTT messages entering local SAN 65 destined for an initiator device such as initiator 35, and applies the appropriate similar flow control measures as discussed above depending on the buffer resources available to switch 50.

[57] In one embodiment, the mechanism for flow control for write operations differs depending on whether the Fast-Write mechanism (discussed below) is enabled or not. For the following, "Transparent Mode" refers to when Fast Write is disabled, since RTT messages are transparent between target and initiator, and "Non-Transparent Mode" refers to when Fast Write is enabled, since in this mode (as will be described) the initiator-side switch issues RTT messages on behalf of the target, in order to optimize SCSI performance in a high-latency environment.

[58] *Flow Control of Non-Transparent Write Operations (Fast Write Enabled)*

[59] Figure 4 illustrates Congestion Manager flow control using the Fast Write feature according to one embodiment of the present invention. As discussed below, the Fast Write

feature advantageously provides the capability to minimize the effects of latency in long-distance SCSI data transfers. When using the Fast Write capability, switches issue SCSI Ready-to-Transfer (RTT) messages to the initiator on behalf of the target, in order to enhance performance by causing the initiator to "fill the pipe" with write data. Ready-to-Transfer messages are thus not transparent from target to initiator; RTT messages sent by the target are not the same as those received by the initiator.

[60] As shown in Figure 4, for example, when Fast Write is enabled, the initiator-side switch 150 would receive a write command from initiator 135 destined for target 145 and would normally immediately respond to the initiator with an RTT message requesting the write data for the entire write command. However, this could cause congestion and head-of-line blocking if the switch 150 lacked the resources to receive the entire write data. In one embodiment, with Congestion Manager enabled, the RTT message would not be issued until buffer resources for the entire Write operation, or a threshold amount, e.g., 256KB, whichever is smaller, become available (note that 256KB is an arbitrary buffer memory size, and other instantiations of the present invention may use other threshold trigger levels as required). If the original write operation transfers less than 256KB of data, then the RTT is issued to the initiator when the buffer resources for the exact Write request size become available. If the original write operation transfers more than 256KB of data, then a 256KB RTT message is generated to affect a partial data transfer when 256KB of buffer resources become available. Additional 256KB RTT messages are then generated, as additional buffers become available, until all of the data from the original write operation is transferred.

[61] In one embodiment, a rate limiter module, as described above, is implemented in switch 150 and operates in conjunction with the Fast Write module to add delay between the RTT messages so as to help minimize congestion in switch 150, for example by adding delay between the RTT messages sent to initiator 135 so that the write data sent by initiator responsive to the RTT messages is substantially equal to or less than the data transfer rate of the WAN link.

[62] *Flow Control of Transparent Write Operations (Fast Write Disabled)*

[63] Figure 5 illustrates a flow control process when the Fast Write feature is disabled. When Fast Write is disabled, RTT messages are passed transparently from target to initiator. In this case, the initiator-side switch 150 only needs to monitor RTT messages coming from the target 145 (or switch 140 in the configuration as shown). If the size of the RTT message is greater than the buffer resources switch 150 has available, then switch 150 holds the RTT message until buffer resources become sufficient to receive the entire write data specified by

the RTT message. Any additional RTT messages that arrive at switch 150 from other initiators are similarly held and serviced as buffer resources become available.

[64] In one embodiment, a rate limiter module, as described above, is implemented in switch 150. The rate limiter module operates to add delay between the RTT messages so as to help minimize congestion in switch 150, for example by adding delay between the RTT messages sent to initiator 135 so that the write data sent by initiator 135 responsive to the RTT messages is substantially equal to or less than the data transfer rate of the WAN link.

[65] *Fast-Write Software Feature*

[66] In one embodiment, a Fast-Write software feature useful in switches and other network devices is provided. The Fast-Write feature of the present invention significantly improves the performance of write operations between Fibre Channel initiators and targets on a wide area network. The actual improvement is dependent on several factors including Wide Area Network (WAN) Round Trip (RT) Time, available buffer on the target (i.e., size of the RTT message), number of concurrent SCSI tasks (e.g., I/O operations) supported by the application, number of concurrent RTT's supported by the target, and the application I/O size. As an example, test results comparing switches enabled with the Fast Write feature of the present invention to Fast-Write disabled switches have shown over a 10X performance improvement for write operations with a WAN delay of 40ms, 740KB I/O transactions, and 8 concurrent I/Os.

[67] *Example of Write transaction without Fast Write*

[68] Figure 6 shows a standard 1MB write transaction.

[69] In the example of Figure 6, the 1MB write transaction begins with the SCSI initiator 235 initiating a SCSI Write Command for 1MB block of data. The command reaches target 245 in $T = 0.5(RT) + 2(\text{local exchange times})$. Compared to the $0.5(RT)$ component the local exchange time can be considered insignificant. The target 245 in this example has 1KB of available buffer, and it responds with a Ready-To-Transfer (RTT) for 1 KB. The RTT is received in $0.5(RT)$ by the SCSI initiator 235. The initiator 235 transmits 1KB of data that is received by the target in $0.5(RT)$. The target 245 then, in this example, has 512KB of available buffer, and it issues a Ready-To-Transfer for 512KB; which is received by the initiator in $0.5(RT)$. The initiator 235 responds and $0.5(RT)$ later the 512KB block is written to the target 245. The process is repeated for the remaining 511KB, and $0.5(RT)$ later all the data is written to the target 245. The target 245 completes the successful operation by issuing a SCSI Response message (FCP_RSP) which is received by the initiator 235 $0.5(RT)$ later. The entire write operation required $4.0(RT)$ or $n+1(RT)$ where n is the

number of RTT issued by the target. In the example of Figure 6, the RTT are represented by FCP_XFER_RDY.

[70] Fast Write Enabled

[71] When the near-end switch 250 (the switch connected to the SCSI initiator) is
5 configured with the Fast Write feature and with Fast Write enabled, as shown in Figure 7, the
near-end switch 250 responds to the SCSI Write command received from initiator 235 with a
Ready-To-Transfer message on behalf of the target 245. Instead of requesting 1KB block of
data, switch 250 requests that the entire 1MB be transferred. Since there is not necessarily
1MB of available buffer on the far-end target 245, the far-end switch 240 stores or caches the
10 1MB of data. The Write command is forwarded to the target 245 where the target requests
data, in this example 1KB of data, be sent by issuing a Ready-To-Transfer message for 1KB.
The far-end switch 240, now storing the 1MB of data, responds with 1KB of data. The target
245 then requests 512 KB, and the far-end switch 240 responds with 512KB of data; and so
forth for the final 511KB. The entire operation takes $1.0(RT) + 12(\text{local exchanges})$. Again,
15 compared to the round trip time component the local exchanges are considered insignificant,
and the net performance improvement is about 4X in this example.

[72] In one embodiment, the Fast Write module of the present invention is implemented in
software, executing on a device processor or a specialized processor module, and provides a
graphical user interface (GUI). In preferred aspects, code for implementing the Fast Write is
20 written in "C" but could be implemented in any language (e.g., assembly, Pascal, etc). Figure
8 illustrates a sample screen provided by software executing in a computer system, such as a
network management computer system, communicably coupled to the network. The GUI
screen shown in Figure 8 is that of a port configuration menu that allows a user, e.g., network
administrator, to selectively optimize WAN throughput using the Fast Write capabilities of
25 the present invention.

[73] It should be appreciated that the Fast Write module of the present invention can be
implemented in a specialized circuit assembly such as an FPGA or ASIC module including
memory. It should also be appreciated that the entire congestion management functionality,
including rate limiter features, although preferably implemented in software, can also be
30 implemented partially or completely in an FPGA or ASIC module.

[74] *Fast Write Conclusion*

[75] The Fast-Write feature of the present invention advantageously (and significantly) improves the performance of write operations from SCSI initiators to SCSI targets.

Improvements will be most significant for WAN links with significant delay, for small write operations, or when there are few concurrent I/Os. While the example above suggests a 4X performance improvement for write operations, improvements greater than 10X have been measured for a WAN link with a RT delay of 70 ms.

[76] In certain preferred aspects, the Congestion Manager, rate limiter and Fast Write modules of the present invention are implemented in network switch devices. However, it should be understood that Congestion Manager and Fast Write modules as described herein can be implemented in any of a variety of other network devices, such as routers, controller cards, gateways, bridges, storage devices, etc. U.S. Patent No. 6,400,730, which is hereby incorporated by reference in its entirety discloses useful network devices, including switch devices, in which modules of the present invention may be implemented.

[77] While the invention has been described by way of example and in terms of the specific embodiments, it is to be understood that the invention is not limited to the disclosed embodiments. To the contrary, it is intended to cover various modifications and similar arrangements as would be apparent to those skilled in the art. Therefore, the scope of the appended claims should be accorded the broadest interpretation so as to encompass all such modifications and similar arrangements.

WHAT IS CLAIMED IS:

- 1 1. A method of reducing network congestion, comprising:
2 receiving a message by a network device coupling a high speed network link
3 with a low speed or TCP/IP based network link, said network device having a buffer
4 memory, said message being sent from a requesting device to a destination device requesting
5 that data be sent over said low speed or TCP/IP based link from the destination device to the
6 requesting device;
7 determining whether the buffer memory has sufficient space to buffer the
8 amount of data identified by the message request; and
9 if so, transferring the message to the destination device and buffering the
10 requested data received from the destination device in response to the message, wherein the
11 requested data is sent over the low speed or TCP/IP based link destined for the requesting
12 device;
13 if not, holding the message until the buffer has sufficient space.
- 1 2. The method of claim 1, wherein the message is a SCSI read request
2 operation identifying data to be read from a target storage device.
- 1 3. The method of claim 1, wherein the network device is a switch device.
- 1 4. The method of claim 1, wherein the message is a SCSI ready-to-
2 transfer (RTT) message responsive to a write request.
- 1 5. The method of claim 4, wherein the requesting device is a storage
2 device responding to the write request and wherein the destination device is a device that sent
3 the write request.
- 1 6. A method of reducing network congestion, comprising:
2 monitoring operation requests received by a network device coupling one or
3 more high speed network links with a low speed network link or TCP/IP based network link,
4 said network device having a buffer memory, said requests being sent between requesting
5 devices and destination devices identifying data to be sent over said low speed or TCP/IP
6 based link;
7 for each received operation request,

8 determining whether the buffer memory has sufficient space to buffer
9 the amount of data identified by the request; and
10 if so, transferring the operation request to the destination device and
11 buffering the identified data received from the destination device, wherein the
12 requested data is sent over the low speed or TCP/IP based link destined for the
13 requesting device;
14 if not, holding the operation request until the buffer has sufficient
15 space.

1 7. The method of claim 6, wherein each operation request includes one of
2 a SCSI read request identifying data to be read from a target storage device and a SCSI
3 Ready-To-Transfer request identifying data to be written to a target storage device.

1 8. The method of claim 6, further including applying a fairness algorithm
2 to determine priority when multiple operation requests are held due to insufficient buffer
3 space.

1 9. A method of enhancing write performance in a network including first
2 and second switch devices coupled over a high latency network link, wherein the first switch
3 device is coupled to an initiator device over a first high speed network link, and wherein the
4 second switch device is coupled to a target device over a second high speed network link, the
5 method comprising:

6 automatically responding to a write request received by the first switch from
7 the initiator with a first ready-to-transfer (RTT) message on behalf of the target device, the
8 first RTT message requesting the write data from the initiator;

9 sending the write request to the target via the second switch device;

10 receiving the write data from the initiator, said write data being sent in
11 response to the first RTT message; and

12 automatically sending the write data from the first switch to the second switch
13 over the high latency network link so that the write data is stored on the second switch
14 device,

15 wherein when the target sends one or more second RTT messages requesting
16 all or a portion of the write data, the second switch is able to respond with the requested
17 amount of the stored write data.

1 10. The method of claim 9, wherein the high latency network link
2 comprises one or a plurality of network links.

1 11. A method of enhancing write performance in a network including first
2 and second switch devices coupled over a first network link, wherein the first switch device is
3 coupled to an initiator device over a second network link, and wherein the second switch
4 device is coupled to a target device over a third network link, wherein the first network link
5 has a slower data transfer rate than both the second and third network links, the method
6 comprising:
7 automatically responding to a write request received by the first switch from
8 the initiator with a transfer message on behalf of the target device, the transfer message
9 requesting the write data from the initiator;
10 sending the write request to the target via the second switch device;
11 receiving the write data from the initiator, said write data being sent in
12 response to the transfer message; and
13 automatically sending the write data from the first switch to the second switch
14 over the first network link so that the write data is stored on the second switch device,
15 wherein when the target sends one or more transfer messages requesting all or
16 a portion of the write data, the second switch is able to immediately respond with the
17 requested amount of the stored write data.

1 12. The method of claim 11, wherein the transfer message is a SCSI ready-
2 to-transfer (RTT) message.

1 13. The method of claim 11, wherein the third network link is a storage
2 area network (SAN).

1 14. The method of claim 11, wherein the third network link includes one
2 or more Fibre Channel links.

1 15. The method of claim 11, wherein the first link is a high latency
2 network link.

1 16. The method of claim 11, wherein the second link is one of a storage
2 area network (SAN) and a local area network (LAN).

1 17. A network switch device, comprising:
2 a first port for coupling to a first network link;
3 a second port for coupling to a second network link, the second network link
4 having a lower data transfer rate than the first network link;
5 a buffer memory; and
6 a congestion management module executing on the switch device, said module
7 being configured to:
8 monitor messages being sent from requesting devices to a destination
9 device requesting that data be sent over said second network link from the destination
10 device to the requesting device; and
11 determine, for each message, whether the buffer memory has sufficient
12 space to buffer the amount of data identified by the message;
13 wherein if it is determined that there is sufficient space, the switch device
14 transfers the message to the destination device and buffers the requested data received from
15 the destination device in response to the message, wherein the requested data is sent over the
16 second network link destined for the requesting device, and if not, the switch device holds the
17 message until the buffer has sufficient space.

1 18. The switch device of claim 17, wherein each message is one of a SCSI
2 formatted read request and a SCSI formatted ready-to-transfer request (RTT).

1 19. The switch device of claim 17, wherein the first network link includes
2 a storage area network, and wherein the second network link includes a wide area network.

1 20. A control module executing on a network switch device, the switch
2 device including a processor, a buffer, a first port for coupling to a low speed or TCP/IP
3 based network link, and one or more second ports for coupling to one or more high speed
4 network links, said module being configured with instructions to:

5 monitor messages being sent from requesting devices to a destination device
6 requesting that data be sent over said low speed or TCP/IP based link from the destination
7 device to the requesting device;

8 determine, for each message, whether the buffer memory has sufficient space
9 to buffer the amount of data identified by the message; and

10 if it is determined that there is sufficient space, to control the switch device to
11 transfer the message to the destination device and buffer the requested data received from the
12 destination device in response to the message, wherein the requested data is sent over the low
13 speed or TCP/IP based link destined for the requesting device; and
14 if not, to control the switch device to hold the message until the buffer has
15 sufficient space.

1 21. The module of claim 20, wherein each message is one of a SCSI
2 formatted read request and a SCSI formatted ready-to-transfer (RTT) request.

1 22. A write enhancement module executing on a network switch device,
2 the switch device including a processor, a buffer, a first port for coupling to a low speed or
3 TCP/IP based network link, and a second port for coupling to a high speed network link, said
4 module being configured with instructions to control the switch to:
5 automatically respond to a write request received by the switch over the high
6 speed network link from an initiator device with a transfer message on behalf of the target
7 device, the transfer message requesting the write data from the initiator;
8 send the write request to the target via a second switch device over the low
9 speed or TCP/IP based network link;
10 receive the write data from the initiator, said write data being sent in response
11 to the transfer message; and
12 automatically send the write data to the second switch over the low speed or
13 TCP/IP based network link so that the write data is stored on the second switch device, such
14 that when the target sends one or more transfer messages requesting all or a portion of the
15 write data, the second switch is able to respond with the requested amount of the stored write
16 data.

1 23. The module of claim 22, wherein the transfer message is a SCSI
2 formatted ready-to-transfer (RTT) message.

1 24. A method of reducing network congestion, comprising:
2 monitoring operation requests received by a network device coupling one or
3 more high speed network links with a low speed or TCP/IP based network link, said network
4 device having a buffer memory, said requests being sent between requesting devices and

5 destination devices identifying data to be sent over said low speed or TCP/IP based network
6 link; and
7 controlling the rate at which the received operation requests are forwarded to
8 the destination devices by the network device based on the amount of data to be returned by
9 the destination devices such that the rate that the requested data is returned to the network
10 device is substantially equal to, or less than, the rate of the low speed or TCP/IP based
11 network link.

1 25. The method of claim 24, wherein controlling includes:
2 determining, for each request, the amount of data to be returned;
3 determining a rate limit based on the data transfer rate of the low speed or
4 TCP/IP based network link; and
5 inserting a time delay between the requests forwarded to the destination
6 devices based on the rate limit.

1 26. The method of claim 24, further including determining an amount of
2 memory currently available in the network device, and temporarily halting the forwarding of
3 the requests until the amount of memory currently available exceeds a threshold amount.

1 27. A control module executing on a network switch device, the switch
2 device including a processor, a buffer, a first port for coupling to a low speed or TCP/IP
3 based network link, and one or more second ports for coupling to one or more high speed
4 network links, said module being configured with instructions to:
5 monitor data request messages being sent between one or more requesting
6 devices and one or more destination devices requesting that data be sent over said low speed
7 or TCP/IP based network link; and
8 control the rate at which the data request messages are forwarded to the
9 destination devices over said high speed network links based on the amount of data to be
10 returned by the destination devices such that the rate that the requested data is returned to the
11 network switch device is less than or substantially equal to the rate of the low speed or
12 TCP/IP based network link.

1 28. The module of claim 27, wherein the instructions to control the rate
2 include instructions to:
3 determine, for each request, the amount of data to be returned;

4 determine a rate limit based on the data transfer rate of the low speed or
5 TCP/IP based network link; and
6 insert a time delay between the requests forwarded to the destination devices
7 based on the rate limit.

1 29. The module of claim 27, wherein the instructions to control the rate
2 include instructions to:
3 forward the data request messages;
4 buffer the requested data received in response to the data request messages;
5 and
6 send the buffered data over the low speed or TCP/IP based network link.

1 30. The method of claim 9, wherein if the target responds with a second
2 RTT message requesting data before the data has been forwarded from the first switch, the
3 second switch immediately forwards the requested data to the target when the requested data
4 is received from the first switch.

1 31. A control module executing on a network switch device, the switch
2 device having a processor, a buffer, a first port for coupling to a low speed or TCP/IP based
3 network link, and one or more second ports for coupling to one or more high speed network
4 links, the module being configured with instructions to:
5 monitor messages being sent between requesting and destination devices
6 requesting that data be sent over the low speed or TCP/IP based link; and for each message:
7 implement a rate limiting function to determine if a message should be
8 forwarded to the destination device,
9 control the switch device to transfer the message to the destination device and
10 buffer the requested data received from the destination device in response to the transferred
11 message if it is determined that a rate limit is not violated, wherein the requested data is sent
12 over the low speed or TCP/IP link destined for the requesting device; and
13 control the switch device to hold the message until the rate limit is no longer
14 violated if sending the message would violate the rate limit.